



Contribution ID: 101

Type: Talk

Evolution of the CERN Tape Archive scheduling system

Tuesday 22 October 2024 13:30 (18 minutes)

The CERN Tape Archive (CTA) scheduling system implements the workflow and lifecycle of Archive, Retrieve and Repack requests. The transient metadata for queued requests is stored in the Scheduler backend store (Scheduler DB). In our previous work, we presented the CTA Scheduler together with an objectstore-based implementation of the Scheduler DB. Now with four years of experience in production, the strengths and limitations of this implementation are better understood. While the objectstore-based implementation is highly efficient for FIFO queueing operations (archive/retrieve), non-FIFO operations (delete, priority queues) require some workarounds. The objectstore backend implementation imposes constraints on how the CTA Scheduler code can be modified and is an additional software dependency and technology for developers to learn. This paper discusses an alternate Scheduler DB implementation, based on relational database technology. We include a status report and roadmap.

Primary author: Dr GUENTHER, Jaroslav (CERN)

Co-authors: SMITH, David (CERN); AFONSO, Joao (CERN); LEDUC, Julien (CERN); SKOVOLA, Konstantina; DAVIS, Michael (CERN); BUGEL, Niels Alexander; OLIVER CORTES, Pablo (CERN); BACHMANN, Richard (CERN); ALFAGEME PEREZ, Sergio; BAHYL, Vladimir (CERN)

Presenter: Dr GUENTHER, Jaroslav (CERN)

Session Classification: Parallel (Track 1)

Track Classification: Track 1 - Data and Metadata Organization, Management and Access